

Refine Search

Search Results -

Terms	Documents
L11 and L10	0

Database:

US Pre-Grant Publication Full-Text Database
US Patents Full-Text Database
US OCR Full-Text Database
EPO Abstracts Database
JPO Abstracts Database
Derwent World Patents Index
IBM Technical Disclosure Bulletins

Search:

L12	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
-----	--------------------------	-------------------------------------	--------------------------

Refine Search

Recall Text

Clear

Interrupt

Search History

DATE: Tuesday, April 19, 2005 [Printable Copy](#) [Create Case](#)

Set Name Query

Hit Count **Set Name**
result set

side by side

DB=USOC,EPAB,JPAB,DWPI,TDBD; PLUR=YES; OP=ADJ

<u>L12</u>	L11 and l10	0	<u>L12</u>
<u>L11</u>	l9 with ((position or location) near4 (reserv\$3 or secur\$3))	71	<u>L11</u>
<u>L10</u>	(mirror\$3 or backup or redundant) near4 (storage or memory)	9890	<u>L10</u>
<u>L9</u>	(memory or storage) near4 control\$4	208369	<u>L9</u>

DB=PGPB,USPT; PLUR=YES; OP=ADJ

<u>L8</u>	L7 and l6	0	<u>L8</u>
<u>L7</u>	l1 with ((position or location) near4 (reserv\$3 or secur\$3))	234	<u>L7</u>
<u>L6</u>	L5 and l4	203	<u>L6</u>
<u>L5</u>	L1 with l3	2526	<u>L5</u>
<u>L4</u>	L1 and (point-to-point or (point adj2 point))	23573	<u>L4</u>
<u>L3</u>	(mirror\$3 or backup or redundant) near4 (storage or memory)	20609	<u>L3</u>
<u>L2</u>	(FCAL or fibre channel arbitrated loop) near8 l1	43	<u>L2</u>
<u>L1</u>	(memory or storage) near4 control\$4	276907	<u>L1</u>

END OF SEARCH HISTORY

[Search Results](#)[BROWSE](#)[SEARCH](#)[IEEE XPLORE GUIDE](#)

Results for "(network storage controller<in>metadata) <and> (point-to-point<in>metadata) <and> (reserve position<in>metadata)"

Your search matched 0 of 1150196 documents.

A maximum of 100 results are displayed, 25 to a page, sorted by **Relevance** in **Descending** order.[» View Session History](#)[» New Search](#)**Modify Search**[» Key](#)

IEEE JNL IEEE Journal or Magazine

 Check to search only within this results set

IEE JNL IEE Journal or Magazine

Display Format: Citation Citation & Abstract

IEEE CNF IEEE Conference Proceeding

IEE CNF IEE Conference Proceeding

IEEE STD IEEE Standard

No results were found.

Please edit your search criteria and try again. Refer to the Help pages if you need assistance revising your search.

[Help](#) [Contact Us](#) [Privacy](#)

© Copyright 2005 IE



Terms used

[network storage controller](#) and [point to point](#) and [mirror](#) or [backup](#) and [reserve position](#)Found 20,191 of
153,034Sort results
by relevance Save results to a Binder[Try an Advanced Search](#)Display
results expanded form Search TipsTry this search in [The ACM Guide](#) Open results in a new window

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale

1 [Highly available systems for database applications](#)

Won Kim

March 1984 **ACM Computing Surveys (CSUR)**, Volume 16 Issue 1Full text available: [pdf\(2.43 MB\)](#)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

As users entrust more and more of their applications to computer systems, the need for systems that are continuously operational (24 hours per day) has become even greater. This paper presents a survey and analysis of representative architectures and techniques that have been developed for constructing highly available systems for database applications. It then proposes a design of a distributed software subsystem that can serve as a unified framework for constructing database applica ...

2 [Pen computing: a technology overview and a vision](#)

André Meyer

July 1995 **ACM SIGCHI Bulletin**, Volume 27 Issue 3Full text available: [pdf\(5.14 MB\)](#)Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

This work gives an overview of a new technology that is attracting growing interest in public as well as in the computer industry itself. The visible difference from other technologies is in the use of a pen or pencil as the primary means of interaction between a user and a machine, picking up the familiar pen and paper interface metaphor. From this follows a set of consequences that will be analyzed and put into context with other emerging technologies and visions. Starting with a short historic ...

3 [The TickerTAIP parallel RAID architecture](#)

Pei Cao, Swee Boon Lin, Shivakumar Venkataraman, John Wilkes

August 1994 **ACM Transactions on Computer Systems (TOCS)**, Volume 12 Issue 3Full text available: [pdf\(2.04 MB\)](#)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Traditional disk arrays have a centralized architecture, with a single controller through which all requests flow. Such a controller is a single point of failure, and its performance limits the maximum number of disks to which the array can scale. We describe TickerTAIP, a parallel architecture for disk arrays that distributes the controller functions across several loosely coupled processors. The result is better scalability, fault tolerance, and flexibility. This article present ...

Keywords: RAID disk array, decentralized parity calculation, disk scheduling, distributed controller, fault tolerance, parallel controller, performance simulation

4 [Local networks](#)

William Stallings

The rapidly evolving field of local network technology has produced a steady stream of local network products in recent years. The IEEE 802 standards that are now taking shape, because of their complexity, do little to narrow the range of alternative technical approaches and at the same time encourage more vendors into the field. The purpose of this paper is to present a systematic, organized overview of the alternative architectures for and design approaches to local networks.

...

5 Compiler transformations for high-performance computing

David F. Bacon, Susan L. Graham, Oliver J. Sharp

December 1994 **ACM Computing Surveys (CSUR)**, Volume 26 Issue 4

In the last three decades a large number of compiler transformations for optimizing programs have been implemented. Most optimizations for uniprocessors reduce the number of instructions executed by the program using transformations based on the analysis of scalar quantities and data-flow techniques. In contrast, optimizations for high-performance superscalar, vector, and parallel processors maximize parallelism and memory locality with transformations that rely on tracking the properties o ...

Keywords: compilation, dependence analysis, locality, multiprocessors, optimization, parallelism, superscalar processors, vectorization

6 The TickerTAIP parallel RAID architecture

Pei Cao, Swee Boon Lim, Shivakumar Venkataraman, John Wilkes

May 1993 **ACM SIGARCH Computer Architecture News , Proceedings of the 20th annual international symposium on Computer architecture**, Volume 21 Issue 2

Traditional disk arrays have a centralized architecture, with a single controller through which all requests flow. Such a controller is a single point of failure, and its performance limits the maximum size that the array can grow to. We describe here TickerTAIP, a parallel architecture for disk arrays that distributed the controller functions across several loosely-coupled processors. The result is better scalability, fault tolerance, and flexibility. This paper presents the Tic ...

7 A NonStop kernel

Joel F. Bartlett

December 1981 **Proceedings of the eighth ACM symposium on Operating systems principles**

The Tandem NonStop System is a fault-tolerant [1], expandable, and distributed computer system designed expressly for online transaction processing. This paper describes the key primitives of the kernel of the operating system. The first section describes the basic hardware building blocks and introduces their software analogs: processes and messages. Using these primitives, a mechanism that allows fault-tolerant resource access, the process-pair, is described. The paper concludes with some ...

8 RAID: high-performance, reliable secondary storage

Peter M. Chen, Edward K. Lee, Garth A. Gibson, Randy H. Katz, David A. Patterson

June 1994 **ACM Computing Surveys (CSUR)**, Volume 26 Issue 2

Disk arrays were proposed in the 1980s as a way to use parallelism between multiple disks to improve aggregate I/O performance. Today they appear in the product lines of most major computer manufacturers. This article gives a comprehensive overview of disk arrays and provides a framework in which to organize current and future work. First, the article introduces disk technology and reviews the driving forces that have popularized disk arrays: performance and reliability. It discusses the tw ...

Keywords: RAID, disk array, parallel I/O, redundancy, storage, striping

9 The TWA reservation system

David Gifford, Alfred Spector

July 1984 **Communications of the ACM**, Volume 27 Issue 7

Full text available:  pdf(2.35 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Where can you find a solid, forthright overview of the computer systems and management behind airline reservations? NASA's space shuttle? Or any of the multitude of other large computer systems that support important projects or national activities? It's hard, sometimes impossible: partly because the people who worked on such systems often do not have the time to write about their experiences: and partly because many professional journalists who interview these people do not have the techni ...

Keywords: ACP, PARS, airline reservation system

10 Recovery in the Calypso file system

Murthy Devarakonda, Bill Kish, Ajay Mohindra

August 1996 **ACM Transactions on Computer Systems (TOCS)**, Volume 14 Issue 3

Full text available:  pdf(318.88 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

This article presents the deign and implementation of the recovery scheme in Calypso. Calypso is a cluster-optimized, distributed file system for UNIX clusters. As in Sprite and AFS, Calypso servers are stateful and scale well to a large number of clients. The recovery scheme in Calypso is nondisruptive, meaning that open files remain open, client modified data are saved, and in-flight operations are properly handled across server recover. The scheme uses distributed state amount the client ...

Keywords: Calypso, cluster systems, distributed state, state reconstruction

11 The HP AutoRAID hierarchical storage system

J. Wilkes, R. Golding, C. Staelin, T. Sullivan

December 1995 **ACM SIGOPS Operating Systems Review , Proceedings of the fifteenth ACM symposium on Operating systems principles**, Volume 29 Issue 5

Full text available:  pdf(1.60 MB)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

12 An end-to-end approach to globally scalable network storage

Micah Beck, Terry Moore, James S. Plank

August 2002 **ACM SIGCOMM Computer Communication Review , Proceedings of the 2002 conference on Applications, technologies, architectures, and protocols for computer communications**, Volume 32 Issue 4

Full text available:  pdf(286.82 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper discusses the application of end-to-end design principles, which are characteristic of the architecture of the Internet, to network storage. While putting storage into the network fabric may seem to contradict end-to-end arguments, we try to show not only that there is no contradiction, but also that adherence to such an approach is the key to achieving true scalability of shared network storage. After discussing end-to-end arguments with respect to several properties of network stora ...

Keywords: IBP, asynchronous communications, end-to-end design, exNode, internet backplane protocol, logistical networking, network storage, scalability, store and forward network, wide area storage

13 The Zebra striped network file system

John H. Hartman, John K. Ousterhout

August 1995 **ACM Transactions on Computer Systems (TOCS)**, Volume 13 Issue 3

Full text available:  pdf(2.76 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Zebra is a network file system that increases throughput by striping the file data across multiple servers. Rather than striping each file separately, Zebra forms all the new data from each client into a single stream, which it then stripes using an approach similar to a log-structured file system. This provides high performance for writes of small files as well as for reads and writes of large files. Zebra also writes parity information in each stripe in the style of RAID disk arrays; this ...

Keywords: RAID, log-based striping, log-structured file system, parity computation

14 Fast detection of communication patterns in distributed executions

Thomas Kunz, Michiel F. H. Seuren

November 1997 **Proceedings of the 1997 conference of the Centre for Advanced Studies on Collaborative research**

Full text available:  pdf(4.21 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Understanding distributed applications is a tedious and difficult task. Visualizations based on process-time diagrams are often used to obtain a better understanding of the execution of the application. The visualization tool we use is Poet, an event tracer developed at the University of Waterloo. However, these diagrams are often very complex and do not provide the user with the desired overview of the application. In our experience, such tools display repeated occurrences of non-trivial commun ...

15 Illustrative risks to the public in the use of computer systems and related technology

Peter G. Neumann

January 1996 **ACM SIGSOFT Software Engineering Notes**, Volume 21 Issue 1

Full text available:  pdf(2.54 MB)

Additional Information: [full citation](#)

16 Peer-to-peer data trading to preserve information

Brian F. Cooper, Hector Garcia-Molina

April 2002 **ACM Transactions on Information Systems (TOIS)**, Volume 20 Issue 2

Full text available:  pdf(490.65 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Data archiving systems rely on replication to preserve information. This paper discusses how a network of autonomous archiving sites can trade data to achieve the most reliable replication. A series of binary trades among sites produces a peer-to-peer archiving network. Two trading algorithms are examined, one based on trading collections (even if they are different sizes) and another based on trading equal sized blocks of space (which can then store collections). The concept of *deeds* is ...

Keywords: Data replication, digital archiving, digital library, fault tolerance, resource negotiation

17 Distributed operating systems

Andrew S. Tanenbaum, Robbert Van Renesse

December 1985 **ACM Computing Surveys (CSUR)**, Volume 17 Issue 4

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index](#)

Distributed operating systems have many aspects in common with centralized ones, but they also differ in certain ways. This paper is intended as an introduction to distributed operating systems, and especially to current university research about them. After a discussion of what constitutes a distributed operating system and how it is distinguished from a computer network, various key design issues are discussed. Then several examples of current research projects are examined in some detail ...

18 Special issue: AI in engineering 

D. Sriram, R. Joobani

January 1985 **ACM SIGART Bulletin**, Issue 91

Full text available:  pdf(8.79 MB) Additional Information: [full citation](#), [abstract](#)

The papers in this special issue were compiled from responses to the announcement in the July 1984 issue of the SIGART newsletter and notices posted over the ARPAnet. The interest being shown in this area is reflected in the sixty papers received from over six countries. About half the papers were received over the computer network.

19 Xunet 2: lessons from an early wide-area ATM testbed 

Charles R. Kalmanek, Srinivasan Keshav, William T. Marshall, Samuel P. Morgan, Robert C. Restrick

February 1997 **IEEE/ACM Transactions on Networking (TON)**, Volume 5 Issue 1

Full text available:  pdf(231.69 KB) Additional Information: [full citation](#), [references](#), [index terms](#)

Keywords: asynchronous transfer mode, available bit rate, constant bit rate, variable bit rate

20 Summary of the Second International Workshop on Network and Operating System Support for Digital Audio and Video 

Ralf Guido Herrtwich

April 1992 **ACM SIGOPS Operating Systems Review**, Volume 26 Issue 2

Full text available:  pdf(2.58 MB) Additional Information: [full citation](#), [index terms](#)

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2005 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)